REPORT DOCUMENTATION PAGE

1b RESTRICTIVE MARKINGS	
2a. SECURITY CLASSIFICATION AUTHORITY	**3 DISTRIBUTION/AVAILABILITY OF REPORT** Approved for public release: distribution unlimited
2b. DECLASSIFICATION/DOWNGRADING SCHEDULE MAR 02 1993	
4. PERFORMING ORGANIZATION REPORT NUMBER(S)	**5 MONITORING ORGANIZATION REPORT NUMBER(S)** AFOSR-TR- 93 0048

6a NAME OF PERFORMING ORGANIZATION Mass. Institute of Technology	**6b OFFICE SYMBOL** *(If applicable)*	**7a NAME OF MONITORING ORGANIZATION** Air Force Office of Scientific Res. N/L
6c. ADDRESS (City, State, and ZIP Code) Dept. of Brain & Cognitive Sciences 79 Amherst St./E10-120/Cambridge, MA 02139		**7b. ADDRESS (City, State, and ZIP Code)** Bldg. 410 Bolling Air Force Base Washington, DC 20332-6448
8a. NAME OF FUNDING/SPONSORING ORGANIZATION AFOSR	**8b OFFICE SYMBOL** *(If applicable)* NL	**9 PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER** AFOSR 89-0504

8c. ADDRESS (City, State, and ZIP Code) Same as 7b	**10 SOURCE OF FUNDING NUMBERS**			
	PROGRAM ELEMENT NO. 61102F	PROJECT NO. 2313	TASK NO. A9	WORK UNIT ACCESSION NO.

11. TITLE (Include Security Classification)
"Top-Down" Influences on "Bottom-Up" Processing

12. PERSONAL AUTHOR(S)

13a. TYPE OF REPORT Annual	**13b. TIME COVERED** FROM 1/92 TO 12/92	**14. DATE OF REPORT (Year, Month, Day)** 12 Jan 1993	**15. PAGE COUNT** 16

16. SUPPLEMENTARY NOTATION

17.	COSATI CODES		**18 SUBJECT TERMS (Continue on reverse if necessary and identify by block number)**
FIELD	**GROUP**	**SUB-GROUP**	Vision, AI, Cognition, Neurophysiology, Visual Psychophysics, Dynamical Systems

19. ABSTRACT (Continue on reverse if necessary and identify by block number)

Although perception is the subject of extensive study, there has been no formal definition of this state. We offer one, and show how even a rather simple formal conceptualization of a percept entails a considerable amount of machinery. Over the past year or two, several components of the machinery required have been studied. These include (i) the role of especially powerful features, called "Key Features", and (ii) how perceptual categories incorporate world knowledge. There also has been progress in understanding preferences for certain structures, as well as in the dynamics of altering preferences (Chaos in Percepts).

20 DISTRIBUTION/AVAILABILITY OF ABSTRACT ☐ UNCLASSIFIED/UNLIMITED ☒ SAME AS RPT ☐ DTIC USERS	**21. ABSTRACT SECURITY CLASSIFICATION** Unclassified
22a NAME OF RESPONSIBLE INDIVIDUAL Dr. John Tangney	**22b TELEPHONE (Include Area Code)** (202) 767-5021	**22c OFFICE SYMBOL** NL

DD FORM 1473, 84 MAR   83 APR edition may be used until exhausted
All other editions are obsolete   SECURITY CLASSIFICATION OF THIS PAGE

JAN 1993

93 3 1 061

Annual Report
on
Top-Down Influences on
Bottom-Up Processing
(January 1993)

Whitman Richards
MIT E10-120
79 Amherst St.
Cambridge, MA 02139

AFOSR 89-0504

617-253-5776
whit@ai.mit.edu

Over the past year, our research effort may be divided into four areas:

1. Formal frameworks for Percepts and Features.

2. Perceptual Categories and World Knowledge.

3. Experiments related to the above.

4. Studies of Dynamical Systems Behavior (Chaos in Percepts).

1.0 Formal Frameworks

Here we have three main thrusts, one concerned with the logical, formal structure that underlies the act of perception (Richards, Jepson). The second is an analysis of contraints upon useful features, and the third is a proposal for how neural machinery might match the incoming sense data to an internal model (Ullman). The work on features is complete. The other two studies are near completion.

1.1 Logic in Percepts (Richards & Jepson)

This work began about three years ago, when we realized that although many are studying "Perception", there is no formal definition of just what a percept is. Without such a definition, how can we decide whether a particular machine or biological state (or model output) qualifies as a perception? Furthermore, how can we build a true theory of a percept without a clear specification of the kinds of state variables, operations, and "language" that are entailed?

Our first answer to "What Is a Percept?" was to note that perceptions are inductive inferences. When conclusions about a state in the world are drawn from

the sense data, then (fallible) premises must be proposed to complete the inference process. Because these premises are fallible – they are simply intelligent guesses – a partial order can be placed upon possible interpretations of the sense data, given the chosen premises. The order is determined by ranking the premise combination that must be "given up". Within such an order, a percept can then be defined as a maximal node. (This is not equivalent to minimizing the faulted premises.) The key to locating these maximal nodes is to be able to reason about the consistency of the data, given the current state of "top-down" knowledge (Jepson & Richards, 1991). In a recent paper, "Lattice Framework for Integrating Vision Modules", we compare a specialized version of our proposal to several others, such as probabalistic reasoning and Hough transform schemes that are often used to resolve conflicting conclusions reached by different sense modules. (A simple example of such a conflict would be when you view the TV screen: motion information implies the scene is three dimensional, but your binocular system claims the scene is flat.)

Over the past year we have considerably tightened the formal underpinnings of our theory. In addition some major changes have been introduced, the most significant being the use of elemental preference relations, from which a lattice of preference orderings can be built. The theory now details aspects of the perceptual process that previously have been ignored, such as the ability to "project" or "simulate" the effects of parameter variations in the internal model used to explain the data, or making explicit note of the types of decision rules that may be used to choose maximally preferred preference states. This work is nearing completion, with a Tech Report planned by the end of January 1993. At the same time, the manuscript will be sent to the journal *Perception*. This paper provides a formal foundation for the type of research underway, and hence is an important complement to experimental studies. An abbreviated version will be presented at a meeting on Cape Cod during mid-January 1993.

1.2 What Makes a Good Feature? (Jepson & Richards)

Here we specify conditions that must be met if a feature is to be a reliable indicator of a world property. This work is available as AI Memo 1356 and also will appear in *Spatial Vision in Humans and Robots*, Cambridge 1993.

Previously, others had proposed that useful features reflect "non-accidental" or "suspicious" configurations that are especially informative yet typical of the world (such as two parallel lines). Using a Bayesian framework, we show how these intuitions can be made more precise, and in the process show that useful feature-based inferences are highly dependent upon the context in which a feature is observed. For example, an inference supported by a feature at an early stage of processing when the context is relatively open may be nonsense in a more specific context provided by subsequent "higher-level" processing. Therefore, specifiation for a "good feature" requires a specification of the model class that sets the current

context. We propose a general form for the structure of a model class, and use this structure as a basis for enumerating and evalutating appropriate "good features". Our conclusion is that one's cognitive capacities and goals are as important a part of "good features" as are the regularities of the world.

1.3 From Features to Categories (Richards, Feldman & Jepson)

Here we show that features meeting the conditions specified above can provide indices into especially useful categories of visual properties in the world. Then we show that for a given set of elemental concepts, the categories associated with these properties have a natural hierarchical (specialization) structure. This structure provides constraints on the form and type of categories that are inferred when visual objects are classified. Furthermore, the structure provides the opportunity for a "logical regularization" of distorted forms or shapes that are corrupted copies of the categorical prototypes. (See *BMVC'92* paper as well as Section 2.0.)

1.4 Sequencing Streams – A Neural Proposal (Ullman)

At a completely different level, Ullman continues to develop a network heirarchy scheme for how "bottom-up" information comes into register with "top-down" models. The basic process, termed "sequence-seeking", is a search for a sequence of mappings or transformations linking a source and target representation. The search is bidirectional throughout the heirarchy – "bottom-up" as well as "top-down". The novel part of the proposal is that the two searches are performed along two separate, complementary pathways, one ascending, the other descending. When a matching pattern is found, regardless of the level, then a chain of activity linking the source and target is generated, facilitating one particular path in the network. The proposal is largely consistent with what is known about cortical machinery, specifically the interplay between the various visual areas, and hence is a hypothesis about the basic scheme of information processing in the neocortex (and thalamus). Experiments related to this proposal are currently underway – see below.

2.0 Perceptual Categories and World Knowledge (Feldman)

This research constitutes a PhD thesis partially supported under the grant. The abstract follows. Some experimental results are highlighted in Section 3.0.

"What makes a good category? Perceptually natural categories – object classes in which an infinity of distinct forms collapse compellingly into a unary description, such as *triangle*, or *dot on a line* – impose structure onto our perceived world. This thesis investigates the formal composition of simple category models, and the properties that distinguish such categories from arbitrary incoherent sets of unrelated

objects. The goal is a formal characterization of human category inferences, including the rather subtle relationship between a perceiver's existing concepts and entailed inductive hypotheses. A critical issue is the formal relationship between mental models and actual world regularities (i.e. covariation in the world among logically orthogonal properties, or "natural modes"). The main formal structure is a lattice of category models, a relational structure that enumerates the various distinct uniform category models in a model class. The lattice serves as a kind of category hypothesis generator, providing the observer with a closed class of distinct models from which to select, each of which corresponds to a coherent "causal" model of the induced category. A computer program is developed to check the validity of the theory, and to generate the lattice of category models for complex families.

A series of experiments are reported in which subjects were asked to induce simple categories from a very small set of unfamiliar sample objects (either one or three objects), and generate novel examples of the category. The results corroborate the lattice theory, and lobby against a view of categorization as any kind of a statistical summary of environmental frequency distributions. In several conditions, subjects produced a frequency distribution that actually contained a larger number of modes (peaks) than there were objects in the sample set; in another condition, subjects' frequency distributions exhibited a mode in a region of the model space where they never observed any examples; and in another condition, subjects produced a frequency distribution that was distinctly modal in a region of the model space in which distribution they observed was carefully arranged to be perfectly flat. In all these cases, the frequency modes corresponded neatly with nodes on the theoretical category lattice computable from the sample set."

3.0 Experiments

There are three general categories for the experiments underway. The farthest along are those which attempt to dissect the neural machinery (e.g. Configuration Stereopsis, Texture Curvature). Much less advanced, and still largely in the pilot stage, are the experiments that attempt to dissect the machinery underlying a percept. The third experimental area involves a dynamical system analysis, and is presented in a separate section.

3.1 Experiments on Percepts and Categories

In Figure 1 are two illustrations of drawings that lead to multistable percepts. In the left panel, the Necker Cube with handle is typically seen from above as a drawer; however it is also easy to see the array as a cup viewed from below, or as a gasoline can with the handle kitty-corner. In each of these cases, the handle is seen attached to a face of the cube. It is extremely difficult to get the handle to float in space, say in the middle of the cube, or in front at say 0.4 of the perceived distance to
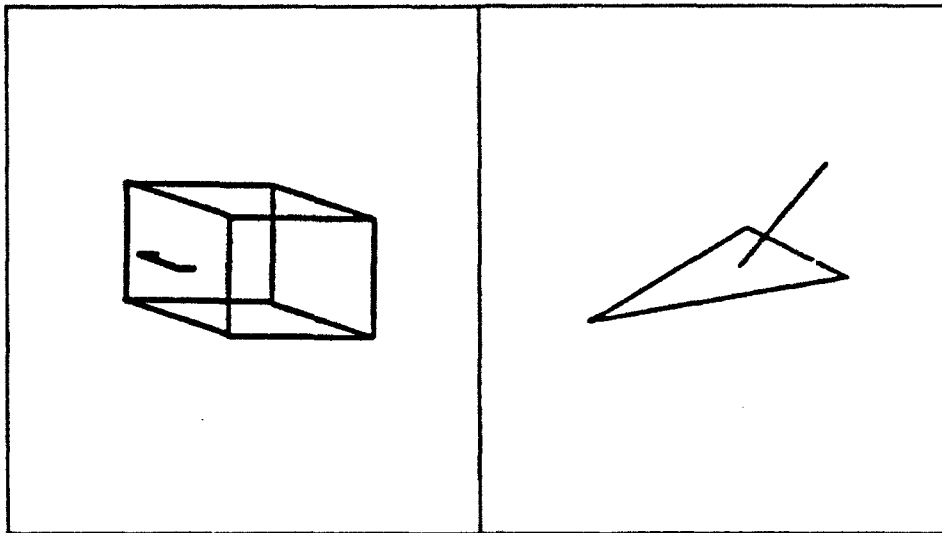
Figure 1 Two examples of drawings with multiple categorical interpretations. For the Necker Cube with handle there are eight common interpretations. For the triangle plus stick there are three. As an example of one preferred state, note that the end of the stick typically lies in the plane of the triangle (just as the feet of the handle lie in the plane of the face). Seeing the stick (or handle) partially penetrating the plane is difficult (see Richards & Jepson, 1992).

the cube. Hence it is obvious that we must have preferred locations for placing the handle along the visual ray. These locations entail a preference for attachment.

The idea behind this set of experiments is to measure the preference strength, or bias, for placing an object, such as the handle, along the visual ray on which it lies. These locations obviously are "set-up" by the structure of the model classes we use to interpret our image data. One question we are studying is whether these states are explored in parallel when the image is analyzed, as suggested by Ullman (1992) in his sequence-seeking model. Or, do we treat each state separately and exclusively, as implied for feature construction? (See also relevant proposals by Koch, 1987; Mumford, 1991, and Carpenter & Grossberg, 1987.)

To fill out the experimental protocol, consider the simple triangle and stick configuration at the right of Figure 1. Most see one of three relations between the stick and the triangle: (a) the stick is upright above the triangle, with its end just touching the plane of the triangle, (b) the stick lies in the plane of the triangle, or (c) the stick lies partly behind the triangle, resting on the side of the triangle, with the left end of the stick in front. (Occasionally people see the stick penetrating the triangle, but this is a rare voluntary initial report.) Elsewhere, we have presented a theoretical analysis for why these three states are chosen (Richards & Jepson, 1992). Here, however, we simply want to prove that any individual has *only* these three possibilities as states in this particular triangle plus stick model, and that there are no other such preferred states. As will be seen shortly, the results will also permit
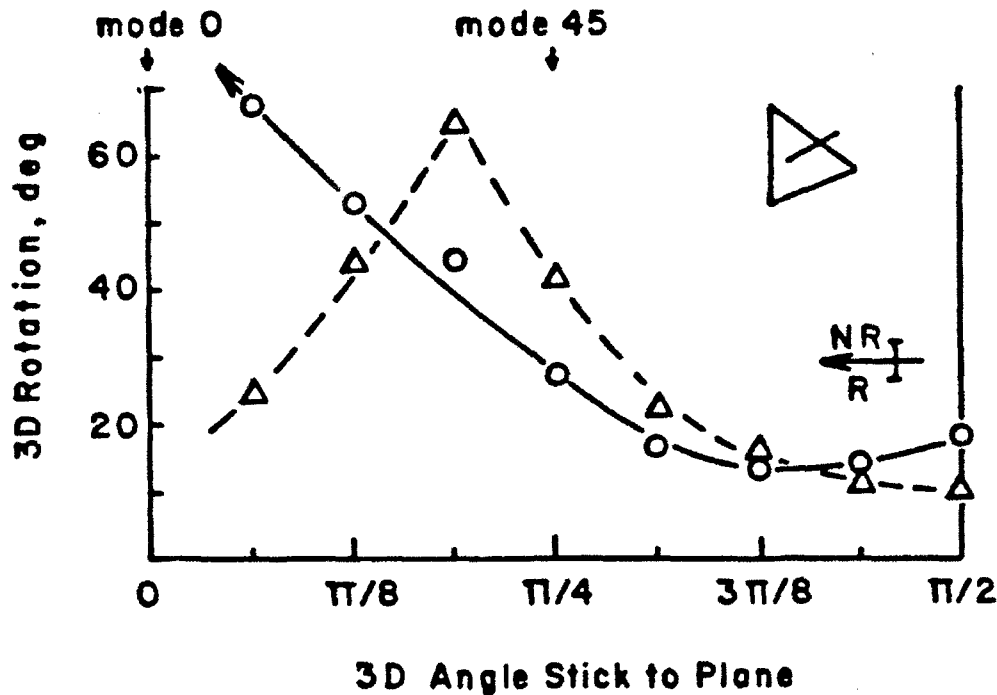
Figure 2 Degree of 3D rotation of triangle and stick needed to offset the bias for a preferred state. States are: "in plane of triangle" and "stick at 45 deg angle to plane of triangle (roughly)". The horizontal arrow at the right indicates amount of rotation that breaks rigidity. The inset shows one view of the configuration.

us to gain some insight as to whether these states are imposed simultaneously on the image analysis, or sequentially, one excluding the other.

One version of the experiment is as follows: generate a 3D representation of the triangle plus stick configuration. Now oscillate this 3D configuration and project the sequence onto the graphics screen, creating a kinetic depth effect.

If the 3D angular rotation is small, then the observer will place the orientation of the stick in his preferred state. As the 3D oscillation increases, however, the correct 3D relation between the stick and the triangle will be noted. Hence the extent of angular rotation of the configuration is a measure of the strength of the preference for a given 3D orientation of the stick to the triangle.

One preliminary set of data are illustrated in Figure 2. The abscissa is the actual 3D angle of the stick to the plane of the triangle, with 0 being the case of the stick in the plane and $\pi/2$ being the case when the stick lies perpendicular to the plane. First, the 3D angular rotation was adjusted until the stick clearly lay off the plane of the triangle (circles). Perhaps not surprisingly, a lot of rotation is required to perceive the stick off the plane when the stick lies near the plane, and little when the stick is perpendicular. Now consider the case when the subject's

bias is to see the stick at roughly a 45° angle to the plane, when the display is static.[1] (See footnote regarding how to estimate this perceived angle for any static configuration.) As seen by the triangular data points for this mode, the greatest amount of 3D rotation lies near 30 degrees (off the predicted mode!), and now more rotation is required in this region than for the planar preference mode (0). Hence, although the configuration remains unchanged, the amount of rotation needed to break a bias depends on the bias present at that moment. Finally, if the judgement is when the stick appears to be articulated (non-rigid relation) or not, then roughly 30° of rotation of the *rigid* array is required regardless of the bias.[2] These results suggest that our preferences play a significant role in the interpretation of the rigid stick-triangle relation as either "stick in the plane of the triangle" or "stick at 45° to the plane".

Regarding whether both states are explored simultaneously during image analysis, we first note that for this subject there is a stronger preference to see the stick in the plane for shallow stick angles [call this state P, and the complementary state O, for off-the-plane], but that the 45° bias (state O) is preferred for intermediate stick angles. What we need to determine is the probability of choosing state O over state P in the early stages of visual processing before a final interpretation is chosen. Our plan is to control the input for state O or P by presenting the configuration stereoscopically in brief flashes. We can then measure the frequency of seeing O or P as function of flash time (and also for the actual 3D angle of the stick to triangle). If both states O and P are initially involved in the analysis, then their relative frequencies should be consistent with Figure 2 as long as the stereo data has not yet been incorporated in the interpretation process. If indeed these relative frequencies remain the same even after it can be shown that the correct 3D slant of the stick has been noted, then this would be evidence that both states O and P are "sent down" together for testing against the data, as Ullman (1992) proposes in his sequencing model. Note, as a bonus, we also will obtain further evidence for distinct preference states using a second psychophysical technique (i.e. stereo vs kinetic depth).

---

[1] This perceived angle can be estimated by first applying Kanade's (1983) skewed-symmetry procedure to the triangle, as if it were isoceles, to determine the surface normal. The maximum likelihood estimate for the 3D angle of the stick can then be shown to be the observed frontal projection of the stick to this normal (or its planar complement).

[2] The fact that this rigid configuration is seen as non-rigid is explained elsewhere (Jepson & Richards, 1992). See also Todd & Bressan, 1990. For shallow stick angles, the articulation is confined to lie in the plane of the triangle. The data should not be examined for consistency – often these kinds of judgements are inconsistent. For example, see Foley, 1972.

3.2 Inherent Structure of Model Classes (with Feldman)

It has become increasingly clear that the perceptual interpretation process relies heavily upon prior knowledge. For example, we readily invoke assumptions about our viewpoint, the expected orientation of a surface, the expected relations between two parts (i.e. attachment preferences discussed earlier), their relative coordinate frames (see Jepson & Richards, 1992; Richards & Jepson, 1992), illuminant position, etc. In many cases, these preferences have an ordering – as in our triangle and stick example. Yet almost nothing is known about how these constraints are organized in our knowledge base. Should rigidity be a special case of articulated motion, or affine motion (as suggested by Ullman & Basri, 1989, or Koenderink & van Doorn, 1990)? Should collinear arrangements be regarded as separate from co-circular, and if not, then just what should be their relation? Should they be separate categories? What about parallel and colinearity?

To address these issues experimentally, we are using a very simple protocal. The subject is given a single exemplar, and asked to draw additional examples. For example, in Figure 3 (top) subjects are shown the drawing in the left panel, then they are asked to draw other members of this category. Typically they will draw more examples of "a dot on a line", allowing the length of the line, its orientation, and the position of the dot to vary. (The lower panel of Figure 3 shows the placements along the line for a collection of subjects.) Why don't the subjects conclude "dot *and* line", placing the dot anywhere, including off the line? Why don't subjects typically place the dot on an extension of the line, or *exactly* at its endpoint? As mentioned in Section 2.0, Feldman (1992) has worked out a theory for this categorical behavior. Again, the idea is that we recognize that a dot on a line required some special attention to its placement (see also the notion of non-accidental features of Binford (1981) and Lowe (1985)). Hence this is a property that must have special significance – in this case in a context of dots and lines thrown out at random. If the sample were a line with the dot *exactly* at the end, then we immediately recognize this case as still more special. In this second case, when subjects are asked to draw more examples from this set, will always place the dot at the line's end – not just anywhere on the line. Hence the "dot at end of line" is more special than the "dot-on-line", which in turn is a special case of "dot *and* line". The cases differ in having degrees of freedom of placement removed (i.e. the codimension of the arrangement goes from 0 to 2). These relations set up a category lattice for "dot-on-line". The subcategories, which entail increasing specialization of placement, each have their distinctive structure. Any example of this structure then indexes to that particular category. (Occasionally the subcategory immediately below will also be included, as dot-at-end of line was in Figure 3. A detailed discussion of this effect and just how one category relates to another can be found elsewhere (Feldman, 1992; Richards et al., 1992).) These set of "dot on line" experiments, as well as a similar set of experiments using two line segments, have been completed and are currently being written up for publication.
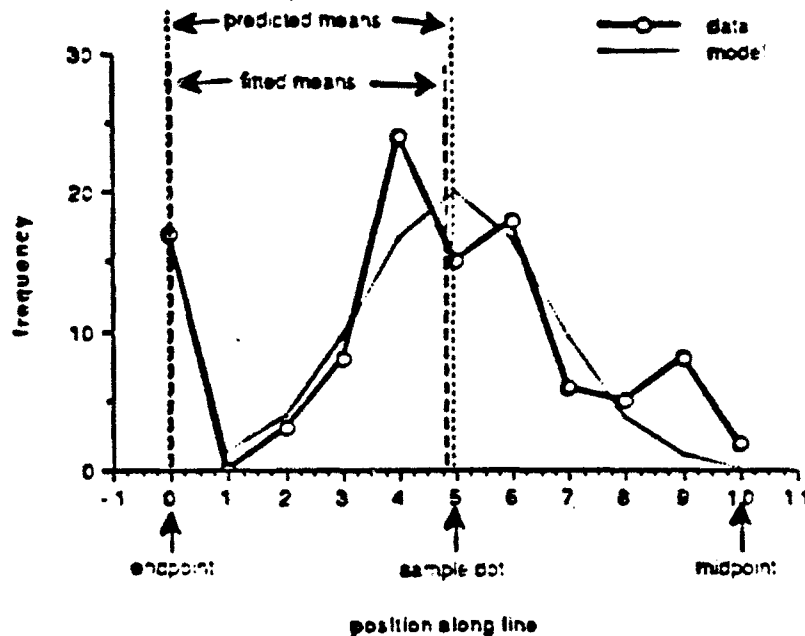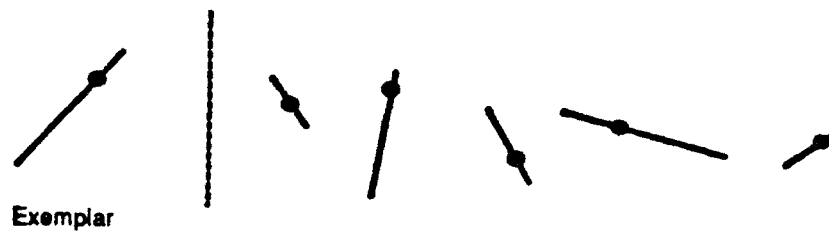
Exemplar

Figure 3 Top: A single "dot on line" is given as an exemplar. Subjects illustrate the category with examples such as the five to the right. Note that orientation, length and dot location were varied. Bottom: distribution of dot locations along the half-line (from Feldman, 1992).

Our next experimental aim is to explore for simple geometrical configurations, the structure of such category lattices and the "features" which index to them. This is a non-trivial problem, because as components and relations are added to create increasingly complex features, the size of the category lattice explodes exponentially. For example, if we have four line segments with the relations parallel equal length, $\pi/2$, touching by end points, then we have 24 possible nodes in the lattice, with the top node being a haphazard arrangement, and the bottom node a square. (As part of his thesis project, Feldman has written a program that automatically generates such lattices – they are too complex to construct correctly by hand.) One reduced version of the 24 node lattice appears on the left panel of Figure 4. (The reduction
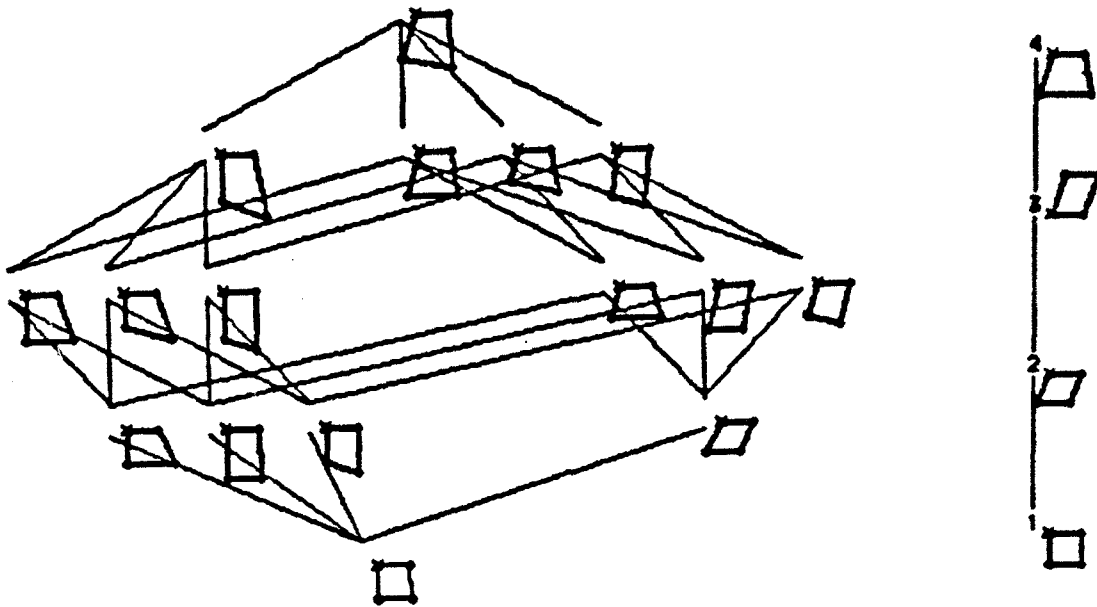
Figure 4 Two reduced versions of the 4-gon lattice, adding a convexity constraint (left) and an "implies causal history" constraint (right). Note that the right lattice is a sublattice of the left. (From Feldman, 1992.)

restricts the 4-gon to be convex.) It is immediately obvious from inspecting this pictorial vision of the 4-gon lattice, that all nodes are not perceptually salient. Typically, when constructing different quadrilateral categories, people will draw a square, rectangle, parallelogram, trapezoid, perhaps a "kite", and a rhombus, such as in the right panel. We are now proceeding to study this 4-gon lattice to determine what further constraints must be placed on the chosen relations in order to obtain a category lattice for quadrilaterals that agrees with our perceptual preferences.

3.3 Configuration Stereopsis (Richards)

This is a completed study on 3D shape that shows how "top-down" information about fixation distance (or shape) modulates angular disparity. Because binocular disparity appears to be computed in V2, this modulation *must* occur early in the visual pathway and hence is potentially accessible to psychophysical probing.

As the distance to an object increases, the angular disparity needed to measure the actual 3D configuration must decrease (reaching zero at the horizon). However, if we take an object, say a cup, and evaluate its 3D shape nearby versus far away, the cup does not appear to flatten, although the disparity signal becomes much

smaller as the distance increases. This suggests a rescaling of disparity with object (or fixation) distance.

Our parametric studies of 3D shape from stereo over a wide range of fixation distances show that indeed, the depth measure associated with a fixed angular disparity changes with fixation distance. The effect is in the direction needed to preserve the shape of 3D configurations as their distance changes, and is roughly two-thirds of what is needed for a full correction. This is evidence for neural signals being modified at or before the extraction of binocular disparity. Hence we have a preliminary "handle" on how a simple case of "bottom-up" information – namely binocular disparity – may incorporate a form of "top-down" knowledge.

This manuscript will be sent off to *Vision Research* near the beginning of February 1993.

3.4 Shading and Stereo (Dawson & Shashua)

Pseudo stereopsis is when the binocular disparities of a surface, such as a face, are reversed but the shading is not. The impression is that the face is "normal" – the nose, for example, still points outward to the viewer.

We have manipulated noses using graphics techniques in order to push them inward, "into the head" so to speak, without altering the shading. No one is able to see these noses "shoved in". Our analysis suggests that this failure of stereopsis is simply due to the shape-from-shading solution "overriding" (in the Percepts Lattice sense) the weak stereo signal created by shaded rather than sharp contours. The effect is not special to faces, and occurs also for "playdo" shapes.

These results need a bit more theoretical work on qualitative shape-from-shading in order to become a complete package. Shashua continues as a post-doc here, and we have set a June 1993 deadline for this project.

3.5 Color Texture (with D.D. Hoffman et al. at Irvine)

Although much work has been devoted to understanding the appearance of homogenous color patches, almost nothing is known about how we represent colored textures. Our approach is to consider the spatial texture pattern as generated by a Markovian process, which "paints" different colors on a surface. The problem, then, is to recover the characteristic parameters of this underlying process.

This problem is almost ideally suited to the formalism described in Observer Mechanics (Bennett, Hoffman & Prakash, 1990), because Markovian kernels lie at the heart of this theory. On the experimental side, we know from earlier work on "Texture Matching" that there will be severe psychophysical restrictions on discriminable patterns, just like in color matching, and expect to find further constraints imposed upon color-texture matches. (Julesz studied this briefly many years ago.)
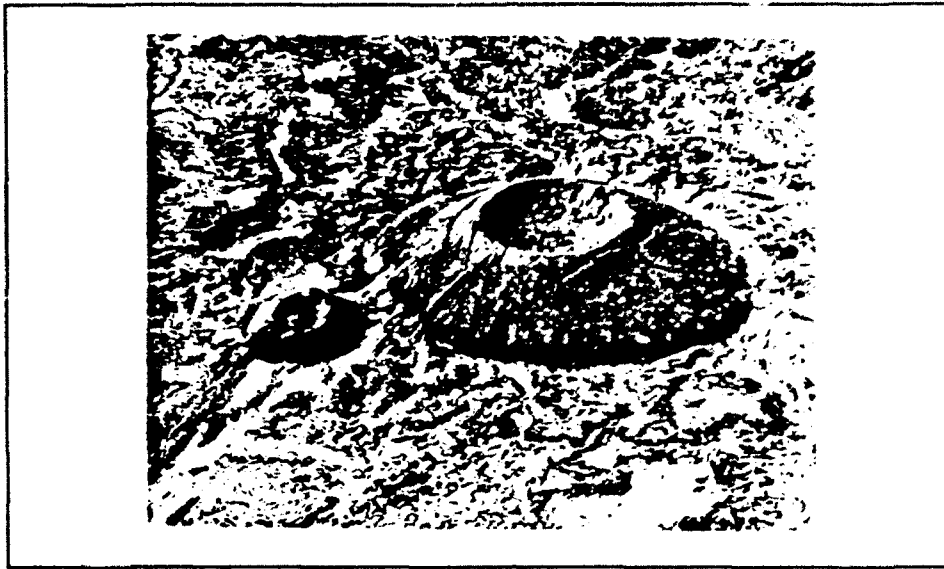
Figure 5 One version of the "crater illusion". (Courtesy of Wide World of Photos, Nov. 1972.)

To date, we have met for three days on this problem at Irvine. We will spend a few more days in January 1993, and another week in the summer of 1993.

3.6 Texture Curvature (with Hugh Wilson)

This study examines curvature discrimination for edges created by texture contours, and includes a model incorporating end-stopped complex cells. The manuscript has appeared in *Jrl. Opt. Soc. A*.

4.0 Dynamical Systems Analysis: Is Perception Chaotic?

The multistability of impoverished visual displays, such as the Necker Cube or the reversible crater illusion illustrated in Figure 5 is well known. What is the dynamics of this switching process? We have analyzed several such perceptual multistabilities, and have found evidence for deterministic chaos in some cases. This work is being prepared for submission to *Science* or *Nature* in January 1993.

Our evidence for deterministic chaos involves a technique that, loosely speaking, measures the fractal dimension of the process that generates the sequence of perceptual transitions. First we measure this time sequence, say obtaining a list of 200 durations. We then compute the average number of intervals $C_p(r)$ whose duration falls within a p-dimensional hypersphere of radius r

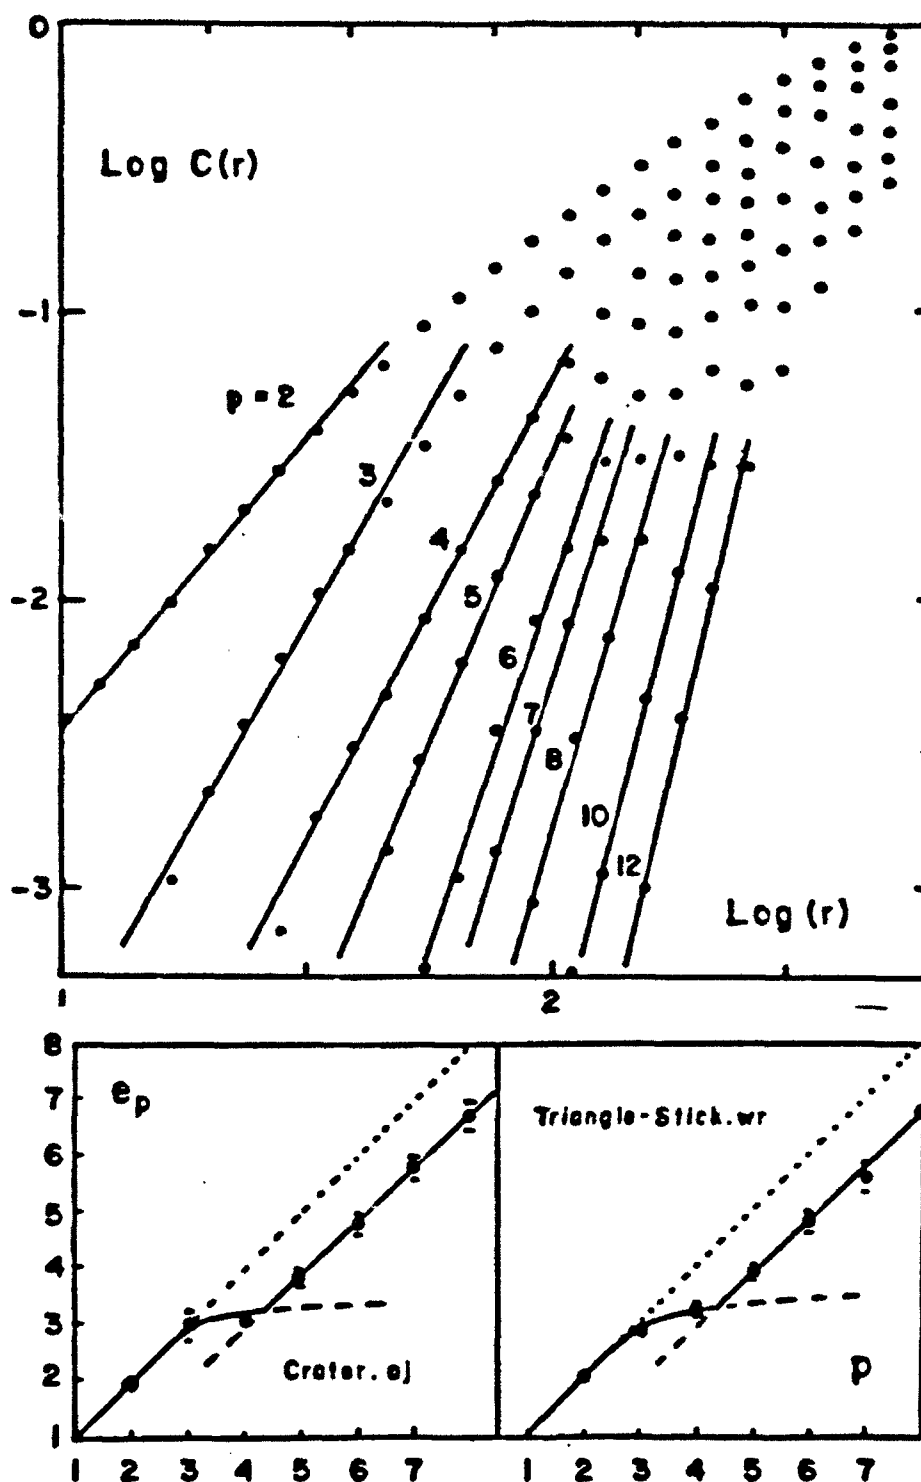$$C_p(r) = \lim \frac{1}{m^2} \sum_{i,j=1 \text{ to } m} H[r - (x_i - x_j)]$$

Figure 6 Top: A plot of $C(r)$ versus r for 200 reversals of the crater illusion.
Bottom left: The exponent e_p taken from the plot above. For a random process
the data points would lie on the dotted line of unit scope. Bottom right: Similar
measurements for 400 reversals of the preferred stick to triangle relations.

where p is the embedding dimension, m is the total number of durations and H is the Heaviside function. (This technique is described clearly in Bergé et al., 1984). For each embedding dimension, an exponent $e_p = \log C(r)/\log(r)$ is calculated and plotted against p. If a random time series is evaluated by this method, $e_p = p$. If a deterministic chaotic series is encountered, such as that for a Henon attraction, then e_p asymptotes at some p_{max} for all $p \geq p_{max}$. Figure 6 (top) illustrates the method. Some preliminary results showing e_p vs p appear in the bottom panels of Figure 6. The lower left panel shows the exponent e_p measured for the reversals in the crater illusion of Figure 5 for subject AJ (200 points). In the right panel, the data were 400 state changes in the position of the stick relative to the triangle discussed earlier (see Figure 1).[3]

Of special interest is the tendency for the perceptual data to asymptote near a value of $e_p = 3.5$, prior to continuing to rise when $p > 4$. (Excluding binocular rivalry, which exhibits behavior typical of a biased random process.) Eye movement patterns taken from a monkey during a search task also show similar behavior. We believe that these results implicate an underlying chaotic process corrupted by noise. To date, we can show that this noise process is not typical of that found in physical devices, such as semiconductors. However, as yet, we do not have an adequate model. The development of such a model is underway.

5.0 Publications (To date)

Bennett, B., Hoffman, D.D & Richards, W. (1991) Reasoning under uncertainty: Lebesgue logic and order logic. *University of California, Irvine, Math. Beh. Sci. Memo MBS 91-08.*

Feldman, J. (1991) Perceptual simplicity and modes of structural generation. *Proc. 13th Ann. Cog. Soc.*, pp. 299-304.

Feldman, J. (1992) Constructing perceptual categories. *Proc. IEEE Computer Soc. Conf. Computer Vision & Pattern Recognition*, pp. 244-250.

Feldman, J., Epstein, D. & Richards, W. (1992) Force dynamics of tempo change in music. *Music Perception*, 10:185-204.

Feldman, J., Jepson, A. & Richards, W. (1992) From features to perceptual categories. *Proc. British Machine Vision Conference, Leeds, 1992*, in press.

Feldman, J., Jepson, A. & Richards, W. (1992) Is perception for real? *Proc. Conference on Cognition & Representation, SUNY Buffalo, April 1992*, in press.

---

[3] Our procedure for estimating e_p is as follows: (1) discard all points where $\sum H = 1 < 10$. (ii) use the max slope for 6 adjacent points, provided that all points lie within 10% of the regression line. (iii) if not, then repeat using 5 points, averaging the two slopes. (iv) if the 10% rule still fails, average the slopes of three lines of four points, etc.

Jepson, A. & Richards, W. (1991) What is a percept? Cognitive Science Paper #43, Center for Cognitive Science, MIT.

Jepson, A. & Richards, W. (1992) A lattice framework for integrating vision modules. *IEEE Trans. Sys., Man. & Cyber.*, **22**: 1087-1096.

Koenderink, J.J. & Richards, W. (1992) Why Is Snow So Bright? *Jrl. Opt. Soc. Am. A*, **9**:645-648.

Moses, Y., Schechtman, G. & Ullman, S. (1990) Self-calibrated collinearity detector. *Biol. Cybern.*, **63**: 463-475.

Richards, W. & Jepson, A. (1992) What makes a good feature? In L. Harris & M. Jenkin (Eds.), *Spatial Vision in Humans and Robots*, Cambridge Univ. Press. Also MIT Center for Biological Information Processing Paper No. 72 (1992) and MIT AI Memo 1356 (1992).

Ullman, S. (1992) Sequence seeking and counter streams. MIT AI Memo 1311.

Wilson, H. & Richards, W. (1992) Curvature Discrimination at Texture Boundaries. *Jrl. Opt. Soc. Am. A*, **9**:1653-1662.

In Preparation:

Logic in percepts (with A. Jepson).

Configuration stereopsis (W. Richards).

Choosing a coordinate frame (with J. Brian Subirana-Vilanova). (See "Figure-ground in visual perception" ARVO 1991 for brief presentation.

Shading and stereo (with B. Dawson & A. Shashua).

Chaos in percepts (M. Somers & H. Wilson).

Talks:

University of Minnesota (May 1989) "'Perception and perceivers".

Harvard University (Nov. 1989) "What's a perception?"

Yale University (May 1990) "What's a percept?"

University of Michigan (June 1990) "What's a percept?"

Cognitive Science Society (July 1990) "Perception, computation and categorization".

Cornell University (June 1991) "What makes a good feature?"

York University (June 1991) "Integrating vision modules".

University of Illinois (Oct. 1991) (1) "What's a percept?, (2) "Choosing coordinate frames".

University of Buffalo (April 1992) "Is perception for real?"

University of Illinois (July 1992) "Inferring perceptual categories". (J. Feldman).

Leeds University (BMVC) (Sept. 1992) "From features to categories".

Brown University (Dec. 1992) "What's a percept?"

6.0 Funds and Personnel

Jacob Feldman was supported in part during 1992 with a portion of the funds originally allocated to S. Ullman.

A no-cost extension of the grant until June 1993 was approved. No funds will be left by the end of the extended grant period.